



ROMANUS: Robust Task Offloading in Modular Multi-Sensor Autonomous Driving Systems

Luke Chen*, Mohanad Odema*, Mohammad Abdullah Al Faruque

Department of Electrical Engineering and Computer Science
University of California, Irvine, California, USA

ABSTRACT

Due to the high performance and safety requirements of self-driving applications, the complexity of modern autonomous driving systems (ADS) has been growing, instigating the need for more sophisticated hardware which could add to the energy footprint of the ADS platform. Addressing this, edge computing is poised to encompass self-driving applications, enabling the compute-intensive autonomy-related tasks to be offloaded for processing at compute-capable edge servers. Nonetheless, the intricate hardware architecture of ADS platforms, in addition to the stringent robustness demands, set forth complications for task offloading which are unique to autonomous driving. Hence, we present ROMANUS, a methodology for robust and efficient task offloading for modular ADS platforms with multi-sensor processing pipelines. Our methodology entails two phases: (i) the introduction of efficient offloading points along the execution path of the involved deep learning models, and (ii) the implementation of a runtime solution based on Deep Reinforcement Learning to adapt the operating mode according to variations in the perceived road scene complexity, network connectivity, and server load. Experiments on the object detection use case demonstrated that our approach is 14.99% more energy-efficient than pure local execution while achieving a 77.06% reduction in risky behavior from a robust-agnostic offloading baseline.

KEYWORDS

Autonomous Driving Systems, Multi-sensor, Sensor Fusion, Vehicular Edge Computing, Robustness, Task Offloading, Object Detection

1 INTRODUCTION

Because erroneous or delayed responses in self-driving applications can compromise road safety, equipment, and/or the lives of the passengers themselves, Autonomous Driving Systems (ADS) are required to achieve outstanding performances on core driving tasks, such as perception and localization. Consequently, ADS platforms are designed today to run highly-sophisticated algorithms on intricate hardware architectures to realize such desired levels of performance while being robust to any adverse driving contexts. For that, modern ADS platforms have adopted a multi-modal processing approach for constructing ensemble perspectives of driving scenes using a diverse set of sensory inputs, as in how the Tesla Autopilot systems possess 8 cameras and 12 ultrasonic sensors [3].

Such a multi-sensor approach leads to the generation of an enormous volume of high-dimensional data that requires tremendous resources for real-time processing, further adding to the power demands of the entire system. Addressing this, a heterogeneous collection of hardware components, as in Application-Specific Integrated Circuits (ASICs) and GPUs, are commonly integrated onto ADS platforms to balance performance demands and power efficiency [13]. Still, hardware advancements are met with growing algorithmic complexity and the requirement for supporting new features, leading the power footprint to remain relatively high. For instance, if we compare two generations of ADS platforms: the Nvidia Drive PX2, which was used by Tesla and Audi Q7 for their autopilot programs [1, 2], against its successor, the Nvidia Drive AGX Orin [4], we find that performance efficiency aside, the baseline power demands increased from 250 W to 800 W, which in theory can have adverse effects on both the thermal comfort of the passengers and the vehicle's driving range [13].

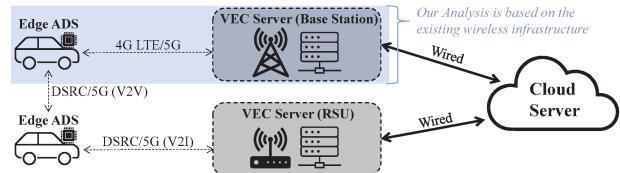


Figure 1: Vehicular Edge Computing (VEC) Architecture

Given how the bulk of processing in ADS modules is largely dominated by deep neural networks (DNNs), compression techniques, e.g., quantization and pruning, have been considered to reduce the modules' complexity, and in turn, their resource requirements [12]. However, experienced performance degradation poses a concern with regards to adopting such techniques' for this class of critical applications. Alternatively, recent research efforts have targeted exploiting the emerging edge computing paradigm for autonomous driving and other vehicular services, in which cloud computing capabilities are brought to the edge of the network through edge servers deployed close to the edge devices, enabling offloading of cumbersome processing burdens to these edge servers for better resource management [14]. In the context of vehicular applications, the edge devices are known as connected vehicles (CVs) and the paradigm is further specified as vehicular edge computing (VEC), or vehicular fog computing (VFC) in other cases [5].

Figure 1 illustrates the hierarchical architecture of VEC and its reliance on the wireless infrastructure, where the higher bandwidths and ultra-low latencies promised by the forthcoming 5G and Dedicated Short Range Communications (DSRC) technologies are to be instrumental in VEC's wide-scale adoption [14]. In this regard, VEC computing servers are expected to be deployed either at Road Side

* Both authors contributed equally to this research.

This work was partially supported by the National Science Foundation (NSF) under award CCF-2140154.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICCAD '22, October 30–November 3, 2022, San Diego, CA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9217-4/22/10.

<https://doi.org/10.1145/3508352.3549356>

Units (RSUs) – as part of the V2X paradigm – or cellular base stations, where recent works have proposed to optimize the offloading process to minimize the overall latency and the energy consumed by the ADS [5, 16]. Still, we find that current approaches in the literature are lacking in the following departments:

- The driving context – crucial to robustness – is not factored in the offloading decision. Given how a scene’s complexity directly correlates with the risk level, and delayed responses could lead to consequences with different levels of severity
- Adapting the offloading load according to the runtime conditions while accounting for the underlying ADS composition (e.g., concurrent sensor processing pipelines) is overlooked

1.1 Motivational Example

In Figure 2, we show two frames from the Radiate dataset [25], and compare their mean average Precision (mAP) scores on the object detection task. As shown, the left frame instantiates a complex scene with numerous objects of diverse classes, some of them superimposed or obstructed from view leading to relatively low mAP scores. Still, we observe that through fusing the outputs from all sensory pipelines (2 cameras, lidar, and radar), an mAP score of 17.6% is realized, which surpasses the highest score achieved by a standalone sensory pipeline output – 11.7% from the right stereo camera. This alludes to the power of sensor fusion as each sensor can capture its own unique set of features that complement those from other sensors to provide more comprehensive views of the driving scenes. On the flip side, the right frame contains a mere single vehicle that is easily detectable by the standalone camera, achieving a 100% mAP score. From here, we can contemplate the desired behavior when VEC is supported with regards to tuning the operating mode. Specifically, an ADS experiencing plain driving scenes can opt for *offloading* processing loads from a subset of sensory pipelines for resource efficiency, because although delayed server responses could cause some partial outputs to be absent by the execution deadline, partial fusion of the available local outputs would suffice for this time window due to the relative simplicity of the scene. Contrarily, complicated scenes should have all sensor outputs available for fusion to stimulate robustness, which is only achieved during *local execution* mode as the uncertainty of the wireless networks is avoided. This behavior would be learned by our proposed solution as will be detailed in the following sections.

1.2 Novel Contributions

To address the above limitations, we present a methodology for Robust Task Offloading in Modular Multi-Sensor Autonomous Driving Systems, namely ROMANUS. From here, we can summarize the main contributions of this paper as follows:

- We present ROMANUS, a methodology to support efficient and robust offloading for modular ADS platform comprising multiple sensory pipelines with support for sensor fusion.
- As far as our knowledge goes, we are the first to factor the driving context in the offloading decision for the robustness of autonomous driving.
- We integrate optimal offloading points within each sensor processing model to realize a dynamic decision space for the runtime operating modes of the ADS.



Figure 2: Two frames of different complexities showing single camera and late fusion mAP scores and the selected operational modes by our learning-based solution. The bounding boxes indicate the ground truths from the dataset.

- We implement a Deep Reinforcement Learning (DRL) based runtime solution that leverages contextual and temporal correlations in the data to optimize the offloading process for latency, energy, and robustness given concurrent pipelines.
- Experiments on the object detection use-case using a real-world driving dataset and an industry-grade ADS indicate that our approach is 14.99% more energy-efficient than local execution while achieving a 77.06% reduction in risky behavior from a robust-agnostic baseline.

2 RELATED WORKS

Mutli-Sensor Perception: To maximize information extraction from a driving scene, data is collected from a diverse set of sensors, e.g., cameras, lidar, and radar, to promote perception robustness. Mainly, There are two primary schemes for processing these multi-sensor inputs: *early fusion* [24, 30] and *late fusion* [29]. The former combines all sensory features to a single feature at an early point in the ADS pipeline, but is susceptible to sensing noise. Conversely, the latter offers more resilience at the expense of more redundancy across the sensor pipelines. Recent works [17, 18] have also explored the potential of *hybrid fusion* approaches to leverage the best of both worlds, albeit with added implementation complexities. Here, we concentrate our analysis on the standard *late fusion* approach as it is more challenging and understudied from an offloading perspective.

Vehicular Edge Computing (VEC): Numerous research efforts have targeted system-wide resource optimization for VEC through optimal task offloading and scheduling strategies given a variety of servers, vehicles, and tasks [28, 31]. Typically, such strategies are complemented with runtime solutions that can tune the operation according to variations in the deployment environment, such as the network connectivity conditions [6]. Nonetheless, delayed responses from edge servers are not tolerated in autonomous driving application as the safety of the road, vehicles, and passengers [5] can be compromised. Hence, [27] proposed a customized communication protocol for a stable and fast offloading of autonomous driving tasks. Even more so, the authors in [16] proposed a fail-safe routine to re-invoke local computation if responses are delayed beyond a certain threshold to account for the uncertainty of wireless links. Such schemes would be even more convoluted when offloading from multiple concurrent pipelines is considered.

DNN Split Computing: To identify optimal offloading points within DNN architectures, [11, 21] analyzed the expected computation and communication costs for each potential offloading layer. For a considerable number of architectures, either direct raw inputs offloading or pure local execution represented the most efficient option. Therefore, works in [7, 19, 20] proposed to modify a DNN’s structure to include an early optimal offloading layer that shrinks

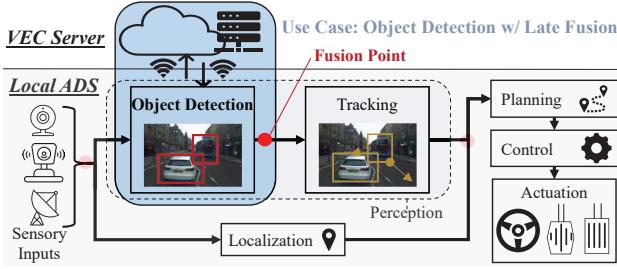


Figure 3: Modular industry-grade ADS. The fusion point following the object detection module is this work’s use-case.

the size of transmissible data, minimizing the costs of both computation and communication. This split-computing concept was applied for end-to-end control in autonomous vehicles [16], and here, we extend its applicability to multi-sensor modular ADS platforms.

3 SYSTEM AND PROBLEM OVERVIEW

3.1 Autonomous Driving System Composition

For perspective, we briefly describe the primary modules that compose a state-of-the-art ADS shown in Figure 3 as follows:

Perception: As the main receptor of the raw sensory data, the perception module is responsible for processing the data over two successive computing blocks. The first is an *object detector* to identify and classify objects of interest, e.g., pedestrians and vehicles, that surround the ego vehicle. A *tracking* module ensues to receive identified objects and associate them with their past movements so as to predict the current movement trajectory.

Localization: Another module taking in the raw inputs is the localization module, whose task is to pinpoint the position of the vehicle at high precision using SLAM/GPS modules.

Planning: Outputs from the perception and localization blocks are fused together onto the same 3D co-ordinate space for further processing by the behavioral and motion planning block. From here, a series of sequential path information can be generated from starting position until the endpoint.

Control: The final block tasked with mapping the information generated from the planning block onto control instructions for the actuators (e.g., driving wheel, brakes, accelerator).

3.2 Problem Formulation

In a modular ADS pipeline, the perception block is the dominant entity affecting the end-to-end performance the most [13], and thus, directing offloading optimizations at this module can lead to substantial efficiency gains across the entire system. Still, sub-optimal operating points can be reached if the following two aspects are not considered properly: (i) the nominal safety considerations of the autonomous driving application, and/or (ii) the structural composition of the ADS modules themselves. For the former, an ADS is required to conclude end-to-end processing under stringent execution time limits to maintain road safety – a 100 ms deadline at the worst [5, 13]. Hence, when VEC is supported, expected additional delays due to wireless channel impairments should be considered as part of the overall end-to-end latencies to determine the best of-flooding decision. Still, additional abrupt delays could threaten the integrity of the self-driving application considering the tightness

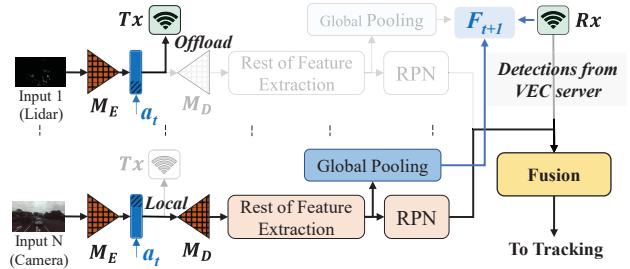


Figure 4: Object Detection with Late Fusion and offloading support. Blue blocks/variables are passed to/from the DRL domain in section 5. Transparent blocks are inactive.

of the execution windows. Whereas for the latter, understanding the underlying structure of an ADS is key to determine the optimal placement of an offloading point that effectively balances the inherent trade-off between communication and computation. For instance, offloading prior to the fusion point can incur a substantial transmission overhead, as opposed to offloading after it which could incur a sizeable computational overhead due to prolonged periods of local processing.

Formally, a module employing late fusion comprises N processing models $\{f_1, f_2, \dots, f_N\}$ for every supported sensor. Thus, for an input vector $X := x_{1:N}$, the fusion block output can be given by:

$$y = \mathcal{H}(f_1(x_1), f_2(x_2), \dots, f_N(x_N)) \quad (1)$$

where \mathcal{H} is the fusion algorithm whose inputs are the N outputs $f_i(x_i) \forall i \in N$. When offloading is supported, the goal is to avoid excessive computational overheads. Thus, each model f_i would incorporate an offloading point to be further defined as:

$$f_i(x_i) = f_i^T(f_i^H(x_i)) \quad (2)$$

where f_i^H and f_i^T are the head and tail parts of the i_{th} model placed prior to and after the offloading point, respectively. The former sub-model is to be deployed locally while the latter is to be replicated across the local and edge server platforms. As server responses could peak due to the wireless channel uncertainty, some model outputs may *not be available* for fusion given the strict execution deadlines. Thus, we obtain instead partial fusion outputs given by:

$$\hat{y} = \mathcal{H}(\mathcal{I}_1 \times f_1(x_1), \mathcal{I}_2 \times f_2(x_2), \dots, \mathcal{I}_N \times f_N(x_N)) \quad (3)$$

where the random variable $\mathcal{I}_i \in \{0, 1\}$ indicates whether $f_i(x_i)$ is available for fusion. Naturally, the lesser number of inputs available the more robustness is compromised. Therefore, given M operating offloading modes, the objective is to identify the mode satisfying:

$$\min_{m \in M} E(X|m), \text{ s.t. } L(X|m) \leq L_{th}, \Delta(y, \hat{y}) \leq C_{th} \quad (4)$$

where $E(\cdot)$ and $L(\cdot)$ are the respective end-to-end energy consumption and execution latency for processing the input vector X given processing mode m . This formulation is regulated by a latency constraint L_{th} for nominal safety, and a robustness constraint C_{th} based on the difference in prediction quality between y and \hat{y} . In the following sections, we present our methodology for solving the optimization objective in (4), which will entail applying DNN structural optimizations and a runtime learning-based approach. We demonstrate our analysis for the use-case of late fusion following the object detection module as illustrated in Figure 3.

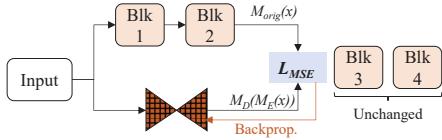


Figure 5: Our ResNet-18 feature extractors undergo Knowledge Distillation to train $\mathcal{M}_E \cdot \mathcal{M}_D$ using the first 2 blocks

4 CONCURRENT PIPELINES DESIGN

Figure 4 depicts the proposed processing domain for object detection with our applied modifications given concurrent DNNs and a late fusion scheme detailed as follows.

4.1 Object Detection with Late Fusion

This scheme entails processing each sensory input separately before aggregating the outputs together through fusion. Specifically, there are two primary computational tasks:

Object Detection: For each sensor, an object detection pipeline is implemented to identify and classify objects in a scene. Initially, each model entails a feature extractor based on a Convolutional Neural Network (CNN), e.g., ResNet-18 here [9], responsible for abstracting raw sensory data into smaller-sized features for the following detection model, e.g., Faster R-CNN network [23], which consists of a regional proposal network (RPN) to suggest regions of interest where objects may exist, a classification stage to categorize the objects within each proposal, and a final post processing stage to convert classified proposals into bounding box predictions.

Fusion: As the outputs from each pipeline are bounding boxes, we can directly fuse them together using Non-Maximum Suppression (NMS) [23] to calculate the intersection over union (IoU) and obtain an estimate on the degree of overlapping between each pair of bounding boxes from the overall set of predictions. If an IoU for a pair of boxes exceeds a predetermined threshold, the bounding box with less confidence score is discarded, and this operation repeats until all possible pairing combinations are covered.

4.2 Implementing DNNs to support Offloading

To avoid the overhead of offloading raw inputs, we scale the optimal offloading point injection technique in [16, 20] to each concurrent pipeline without compromising the overall utility as follows:

Structural Alterations: We alter the structure of the feature extractors (ResNet-18 here) within each pipeline to minimize local computation overhead prior to the offloading point and downsize the transmissible data. Specifically, we substitute a considerable portion from the earliest parts of a DNN with an encoder-decoder like structure of two functional components: (i) an encoder, \mathcal{M}_E , which offers an efficient offloading option at its output through shrinking the input data into a lower-dimensional representation that retains the most relevant of features, i.e., small output sizes translate to low communication overheads., and (ii) a decoder, \mathcal{M}_D , to receive outputs from the encoder and cast them back to higher dimensional representations of dimensions compatible with the remainder of the network. Here, we replace the first two residual blocks from each ResNet-18 with an encoder-decoder structure.

Knowledge Distillation: Next, modified architectures need to be retrained to maintain utility. We apply knowledge distillation

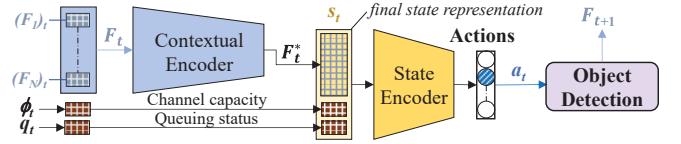


Figure 6: Our hierarchical agent for runtime mode selection

[20] to train \mathcal{M}_E and \mathcal{M}_D through minimizing a loss function, e.g., mean squared error (L_{MSE}), between \mathcal{M}_D outputs and those from the original parts, \mathcal{M}_{orig} . Figure 5 illustrates this for our ResNet-18 with the loss component for a single input x given by:

$$L_{SE} = \|\mathcal{M}_{orig}(x) - \mathcal{M}_D(\mathcal{M}_E(x))\|_2^2 \quad (5)$$

Hence, unaltered DNN components can retain their weight values with only the parameters of the new structure trained to produce the same output values as the originals.

Deployment for Inference: After retraining, the modified architectures would be deployed for each concurrent detection pipeline on the local ADS. Furthermore, each \mathcal{M}_D and its succeeding blocks would be replicated across the VEC servers to enable online autonomous driving services. Thus during runtime, servers can receive outputs from \mathcal{M}_E components, process them, and return predictions, e.g., bounding boxes coordinates, to the ADS platforms. On the ADS, available local and received predictions are fused to provide the final outputs for the following blocks.

5 REINFORCEMENT LEARNING CONTROL

VEC operation is reliant on the surrounding conditions with regards to the wireless channel state and the server load. Hence, we propose a learning-based approach based on deep reinforcement learning (DRL) to adapt the mode of operation so as to maximize performance efficiency while maintaining robustness – which we account for in the offloading decision through leveraging the abstract feature representations already computed within the processing pipelines.

5.1 Hierarchical Agent

As shown in Figure 6, our DRL solution constitutes a hierarchical agent whose main components are as follows:

5.1.1 Contextual Encoder. In order to estimate the complexity of the corresponding scene, we leverage the computed feature set, \mathcal{F}_t , at time window t from the main sensor processing pipelines to guide the decision for the following window $t + 1$, given as $\mathcal{F}_t = \{(\mathcal{F}_1)_t, (\mathcal{F}_2)_t, \dots, (\mathcal{F}_N)_t\}$. The rationale behind using the feature set of the preceding time window is twofold: (i) features do not need to be computed from scratch as they have already been generated within the primary processing pipelines (see the global pooling blocks in Figure 4), and, (ii) the small window size for autonomous driving (≤ 100 ms) means that successive frames share similar driving contexts due to the high spatio-temporal scenic correlations.

Given how \mathcal{F}_t can outweigh other DRL inputs due to its relatively larger size, \mathcal{F}_t needs to be initially encoded into a further lower-dimensional representation. Hence, \mathcal{F}_t is processed through a *contextual encoder* comprising a sequence of fully-connected layers to obtain the final abstraction \mathcal{F}_t^* . In our experiments, \mathcal{F}_t^* was of $256\times$ smaller in size than \mathcal{F}_t .

5.1.2 State Encoder. The next component is the state encoder whose input is the final state representation $s_t = \{\mathcal{F}_t^*, \phi_t, q_t\}$ formed from aggregating the contextual encoder outputs, \mathcal{F}_t^* , the channel capacity ϕ_t , and server queuing delays q_t . Practically, the latter two metrics can be estimated by probing the edge server.

5.1.3 Action Space. Represented by the final fully-connected layer in the *state encoder*, the action space covers the set of all possible modes of operation that can be selected by the DRL at runtime. We define it as $A = \{offload_0, offload_1, offload_2, \dots, offload_{N-1}\}$, where an action $offload_i$ is for choosing the offloading option for i sensory pipelines, with $offload_0$ being pure local execution. In the case that the same DNN structure is shared across all pipelines, only the number of offloading pipelines matter. We do not consider $offload_N$ as a viable option so that the ADS is always guaranteed a new output every time window since at least one pipeline is always processed locally. This way, even under a worst-case scenario when tasks from $N-1$ pipelines are offloaded and results are not received within the time limit, the vehicle can still operate in a safe manner. In practice, we merely need a subset of actions $A^* \subseteq A$, with $\{offload_0, offload_{N-1}\} \subseteq A^*$, where A^* can contain the actions that exhibit notable variability in performance. At runtime, action vector, a_t , is mapped onto the control of each processing pipeline.

5.2 DRL Environment

We detail the emulated DRL training environment for learning a policy π that makes offloading decisions based on the current state.

5.2.1 Training and Reward. Reinforcement learning approaches rely on having a Q function to provide value estimates for each state-action pair so as to select the optimal action $\hat{a} = \arg \max_{a \in A} Q_\pi(\hat{s}, a)$ for each \hat{s} under a learnt policy π . However, estimating state-action pair values in continuous state spaces is challenging, and DRL offers to approximate Q_π by a policy network trained to maximize a reward. With no loss in generality, our DRL employs a Double Deep Q-Network [26] with a compounded reward function \mathcal{R} as:

$$\mathcal{R} = \begin{cases} \mathcal{A}, & \text{if } mAP(y) < mAP_{th} \\ \mathcal{B}, & \text{otherwise} \end{cases} \quad (6)$$

which evaluates to different functions based on a measure of robustness, which we associate here with the degree of uncertainty in the final fused predictions y in (1), determined by the mean Average Precision (mAP) scores for object detectors as in [22]. In brief, our goal is for the agent to realize a policy that deters from offloading actions when prediction confidence is low, which we achieve here through leveraging the contextual information in \mathcal{F}_{t-1} to assess the scene's complexity, and make offloading decisions accordingly with the goal of minimizing prediction uncertainty. Thus, if mAP_{th} is not met, \mathcal{R} evaluates to \mathcal{A} defined as:

$$\mathcal{A} = \begin{cases} 0, & \text{if } \hat{a} == offload_0 \\ \frac{P}{N-i}, & \text{if } \hat{a} == offload_i \quad s.t., i \neq 0, i < N \end{cases} \quad (7)$$

for penalizing the agent whenever an offloading action is selected, with the penalty value being proportionate to the number of offloading pipelines, i , out of N total, with a maximum *negative* penalty of P . Recall that $offload_N \notin A$ as one pipeline always executes locally to ensure at least one output is available irrespective of the wireless

network conditions. On the flip side, when mAP_{th} is satisfied, \mathcal{R} evaluates to \mathcal{B} as follows:

$$\mathcal{B} = \begin{cases} P, & \text{if } L(X|\hat{a}) > L_{th} \\ C, & \text{otherwise} \end{cases} \quad (8)$$

which penalizes the agent by P when its selected action \hat{a} causes the overall execution latency for inputs X , $L(X|\hat{a})$, to exceed the critical execution latency constraint, L_{th} . In other words, this means that the agent is penalized when not all partial outputs are available in time for late fusion. In reality, state-of-the-art ADS platforms are designed to meet the application latency demands, and hence, we set the value of L_{th} to that of *local execution*. Contrarily, when L_{th} is satisfied, \mathcal{R} finally evaluates to C given by:

$$C = \begin{cases} 0, & \text{if } E(X|\hat{a}) == \min(E(X|a) | L(X|a) \leq L_{th}) \\ P, & \text{otherwise} \end{cases} \quad (9)$$

$$\forall a \in A^*, A^* \subseteq A$$

penalizing the agent by P if the energy consumption footprint $E(X|\hat{a})$ from selecting action \hat{a} is not the minimal from amongst those of all viable actions $a \in A^*$ that are projected to meet L_{th} .

5.2.2 Latency and Energy Estimation. In order to evaluate \mathcal{R} for each selected \hat{a} , the end-to-end estimates for energy and latency can be approximated every time window as follows:

$$L = L_{local} + L_{Tx} + L_{server} + L_{Rx} \quad (10)$$

$$E = E_{local} + E_{Tx} + E_{idle} + E_{Rx} \quad (11)$$

where the latency L can be broken down into the respective local, transmission, server, and receiving latencies. Similarly, energy consumption constitutes the same components except for incorporating idling energy as we are only concerned about the ADS energy footprint. From here, the local components are given by:

$$L_{local} = N \times L_{M_E} + (N - i) \times L_{tail} \mid \hat{a} == offload_i \quad (12)$$

in which L_{M_E} and L_{tail} are the respective latencies for executing the encoder M_E and the remaining tail parts of the model, respectively. When the selected action is to offload processing from i processing pipelines (i.e. $\hat{a} == offload_i$), the total local execution latency accounts for processing across the N encoders and the $N - i$ tail models. This additive form represents the most direct approach for modeling local execution. However in reality, the concurrency of pipelines can speed up local execution depending on the available hardware resources at the expense of a larger power consumption footprint, P_{local} . We approximate this trade-off through considering energy for performance evaluation, defining E_{local} as:

$$E_{local} = L_{local} \times P_{local} \quad (13)$$

5.2.3 Channel Estimation. To estimate the communication overheads, we first fit a Rayleigh distribution curve with scale σ to throughput traces Φ collected from the real-world for different wireless technologies, i.e., $\Phi \sim \text{Rayleigh}(\sigma)$. Then, we use the constructed distribution to sample independent and identically distributed (i.i.d.) random variables as the channel capacity ϕ to be used for the training and evaluation processes of the DRL agent where data transmission parameters can be evaluated as:

$$L_{Tx} = \frac{i \times b}{\phi} \mid \hat{a} == offload_i; E_{Tx} = L_{Tx} \times P_{Tx} \quad (14)$$

where b is the transmissible data size from one sensory pipeline while P_{Tx} is the transmission power incurred by the ADS. Similarly, the formulation for the receiving parameters, L_{Rx} and E_{Rx} , can be provided given corresponding estimates for channel capacity and data sizes in the downlink.

5.2.4 Server Queuing. Lastly, we model the server latency L_{server} using queuing delays where we have:

$$q_c = \frac{(1 - \rho)(\rho)^c}{1 - \rho^{C+1}} \quad (15)$$

representing the probability that the offloaded task would encounter c other tasks before it in the server's processing queue, with ρ being the average server load, and C being the queue size. From here, we are able to generate a probability density function (pdf) for values within 0-C from which we can sample queuing positions, and consequently approximate L_{server} .

6 EXPERIMENTS AND RESULTS

6.1 Experimental Setup

6.1.1 Dataset. We use the RADIATE multimodal perception dataset [25] for its diverse driving scenarios and adverse weather conditions such as snow, fog, and rain. The dataset covers 8 object classes with annotations from a Navtech CTS350-X radar, a Velodyne HDL-32e LiDAR, and a ZED stereo camera. The variety of scenes provides a varying degree of difficulty for ADS and enables the robustness assessment. For instance, cameras obstructed by snow offer poor visibility indicating higher difficulty that can cause sub-optimal object detection performance. Here, we implemented 4 object detection DNN pipelines: 2 stereo cameras, radar, and lidar. All inputs are mapped onto the forward-facing perspective for late fusion.

6.1.2 Training and Metrics. As was mentioned in Section 4, the original processing pipelines for each sensing modality comprise a ResNet-18 followed by a Faster R-CNN. These models were trained using a batch size of 1, learning rate of 0.005, and the multi-task loss function in [23] which combines both classification and box regression losses. For the NMS fusion, we use a fusion IoU threshold of 0.4. We employ mAP as our evaluation metric with boxes IoU ≥ 0.5 since it is widely adopted for object detection tasks [8] where the average precision is estimated using the precision and recall values. More details about evaluating these values are in [8, 23].

6.1.3 Hardware and Performance Evaluation. We use the industry-grade Nvidia Drive PX2 Autochaffer as our ADS hardware. The concurrent DNN models are compiled using the TensorRT library becoming inference engines. The local execution power P_{local} is estimated as the difference in the ADS power measurements when processing and idling. For the transmission power P_{Tx} , we follow [16] and evaluate it using the data transfer power models in [10].

6.1.4 Encoder-Decoder Structure. The input frame's resolution for each of the sensory pipelines is 672×376 (≈ 740.25 kB). The encoder, \mathcal{M}_E , comprises 3 layers (2 convolutional and 1 pooling), each with a stride of 2 with only 3 channels at the output. Therefore, when the outputs from \mathcal{M}_E are quantized to 8 bits for offloading [16], the transmissible data size b in equation 14 becomes ≈ 11.57 kB (64 \times less than the input's). The decoder \mathcal{M}_D mimics the structure

presented in [20] to have its output of the same dimensions as that from the original second ResNet-18 block.

6.1.5 DRL Settings. For safety, we always execute the radar pipeline locally [15] and define $A^* = \{offload_0, offload_2, offload_3\}$. We set $P = -2$, $C = 4000$, $\rho = 0.9$, and $mAP_{th} = 0.68$ unless otherwise stated. We set $L_{th} = 68.12$ based on pure local execution latency.

Table 1: Loss and mAP (%) before (orig) and after (dist) integrating $\mathcal{M}_E \cdot \mathcal{M}_D$ across various late fusion combinations.

Sensor	Loss (orig)	Loss (dist)	mAP (orig)	mAP (dist)
2 Cameras	0.15	0.17	67.14	67.14
Radar+Lidar	0.10	0.11	67.14	67.14
Full Fusion	0.13	0.15	71.24	70.38

6.2 Object Detection and Performance

We first assess how the inclusion of \mathcal{M}_E and \mathcal{M}_D impacts the loss and prediction accuracy of object detection. Table 1 shows the changes in these metrics across different late fusion combinations on the Radiate evaluation dataset. As seen, full sensor fusion has the best performance in mAP, asserting how prediction robustness relates to the number of fused outputs. It is also observed that the new DNN structures maintain the same level of performance as their original counterparts, with the highest degradation in mAP from 71.24% to 70.38% experienced by the full fusion case, but still offering a better score than that of the simpler sensor combinations.

Table 2: Hardware Measurements on the Nvidia Drive PX2

DNN	L_{local} (ms)	E_{local} (J)	Memory (MB)
Encoder	3.78	0.03	0.025
1 pipeline	17.03	0.12	80.3
4 pipelines	68.12	0.48	321.2
DRL Agent	0.66	0.005	5.4

Table 2 displays the processing overheads for different DNN components deployed on the PX2 hardware. The encoder \mathcal{M}_E and DRL agent take 3.78 and 0.66 ms, respectively, emphasizing how the decision a_t is obtained before the generation of any transmissible outputs. Moreover, the execution latency for 4 pipelines on the PX2 can add up to 68.12 ms given the same power P_{local} .

6.3 Channel Capacity and Queuing Analysis

In this experiment, we analyze the influence of the experienced channel capacity, ϕ , and queuing delay, q_t , on the optimal action choice when optimizing for energy consumption under the latency constraint L_{th} . To elaborate, we illustrate in Figure 7 parametric sweeps with respect to ϕ given $q_t = 15$ ms. As shown, offloading options are consistently more energy efficient than the pure local option ($offload_0$), but the L_{th} constraint dictates which action should be chosen considering how poor values of ϕ could disqualify some offloading choices. When $\phi > 4$ Mbps, the latency overhead for $offload_2$ does not exceed L_{th} making it the optimal action until $\phi > 7$ Mbps, at which the most energy-efficient option, $offload_3$, becomes valid. Similarly, this analysis is repeated in Figure 8 when sweeping across q_t given $\phi = 8$ Mbps. Naturally, the latency overhead is linearly proportional to q_t under fixed network conditions,

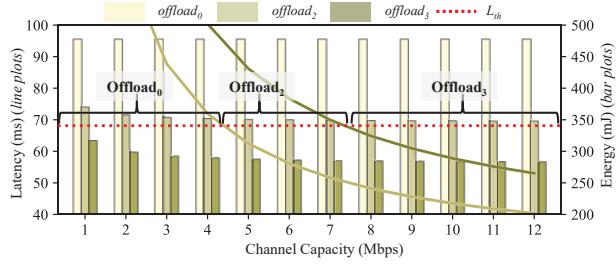


Figure 7: Variation of Latency and Energy Analysis w.r.t. Channel Capacity. Energy as bar charts, Latency as plot lines

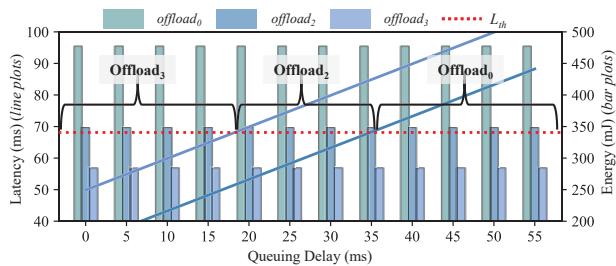


Figure 8: Variation of Latency and Energy Analysis w.r.t. Queuing Delay. Energy as bar charts, Latency as plot lines

demonstrating the influence of server load over the optimal offloading decision. From here, the key takeaway is that based on the wireless infrastructure and VEC server capabilities, the maximum number of concurrent offloading pipelines that meet L_{th} can be determined, and used accordingly to construct the decision space.

Table 3: Robustness analysis at $\rho=0.97$

Policy	mAP _{th}	Average mAP (%)		
		offload ₀	offload ₂	offload ₃
R-agnostic	N/A	64.93	64.50	64.85
	0.50	60.39	67.69	68.58
DRL	0.68	60.68	70.93	72.55
	0.98	61.82	72.49	73.14
Oracle	0.50	49.58	85.34	85.26
	0.68	50.58	93.68	93.62
	0.98	54.84	99.51	99.49

6.4 Robustness Analysis

To analyze the DRLs capacity to make robust decisions, we define 2 baseline policies for comparison: (i) a robustness-agnostic or *R-agnostic* policy that is aware of ϕ and q_t to optimize for energy so long as L_{th} is satisfied, and (ii) an *Oracle* resembling an optimal strategy which in addition to the information available to the *R-agnostic* policy, also possesses the true per-frame mAP estimate apriori, and consequently, the optimal sequence of decisions satisfying the robustness constraint mAP_{th} . All of the mentioned strategies are more energy-efficient than pure local execution, and the mAP_{th} values are set in the experiments to 0.5, 0.68, and 0.98, estimated based on the cumulative distribution of the evaluation dataset such that 30%, 50%, and 70% of the evaluation mAP scores fall under the corresponding thresholds.

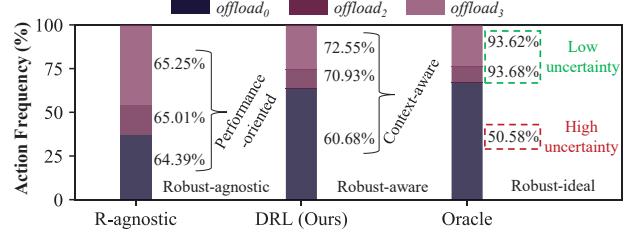


Figure 9: Action selection frequencies (%) breakdown across the 3 policies at $mAP_{th}=0.68$ and $\rho=0.97$. Numbers next to the bars indicate the average experienced mAP (AMAP) for the evaluation dataset inputs mapped to each action.

To evaluate robustness across each policy, we compute the average experienced mAP per action (AMAP) given the action selection frequencies. Mainly, a robust behavior cause frames of high uncertainty ($mAP \leq mAP_{th}$) to be processed locally, implying how the AMAP experienced locally should be low compared to those from the offloading actions. We illustrate this concept in Figure 9 across the 3 policies for $mAP_{th} = 0.68$ and $\rho = 0.97$. As seen, the *R-agnostic* policy only considers performance efficiency for its action selection, and subsequently, its AMAP across $offload_0$, $offload_2$, and $offload_3$, are equivalent with values of 64.39%, 65.01%, and 65.25%, respectively. Conversely, the *Oracle* resembles the ideal embodiment of robustness, assigning high uncertainty frames to $offload_0$, despite performance gains from offloading. In contrast to the *R-agnostic* policy, 66.94% of the *Oracle* policy decisions are $offload_0$ with an AMAP of 50.58%, and an AMAP as high as 93.68% for the remaining offloading decisions. From here, our proposed *DRL* approach strives to learn the *Oracle*'s behavior, through the observed action selection breakdown, with 63.65% of actions belonging to $offload_0$. Moreover, the AMAP for $offload_2$ and $offload_3$ are 70.93% and 72.55%, respectively, which despite outperforming the *R-agnostic* policy, are far from that of the *Oracle*. This is expected considering the *Oracle* policy is the unrealistic ideal behavior with apriori mAP knowledge. We extend this analysis to other thresholds values in Table 3, where we observe that as the robustness constraint becomes smaller, the *DRL* exhibits a behavior closer to the *R-agnostic* and farther from the *Oracle* and vice versa, indicating the *DRL*'s capacity to adapt to various robustness requirements.

Table 4: Action frequency analysis at $mAP_{th} = 0.68$

Policy	ρ	Action Frequency (%)		
		offload ₀	offload ₂	offload ₃
R-agnostic	0.90	11.17	14.73	74.10
	0.97	37.18	17.17	45.65
	0.99	70.34	9.25	20.41
DRL	0.90	51.73	8.83	39.44
	0.97	63.65	10.95	25.39
	0.99	85.14	5.46	9.4
Oracle	0.90	53.19	8.00	38.82
	0.97	66.94	9.17	23.89
	0.99	84.37	5.04	10.60

Furthermore, we vary the server load ρ in Table 4 and show how the action frequency varies for each policy. As ρ increases, the selection of local processing becomes more frequent across all

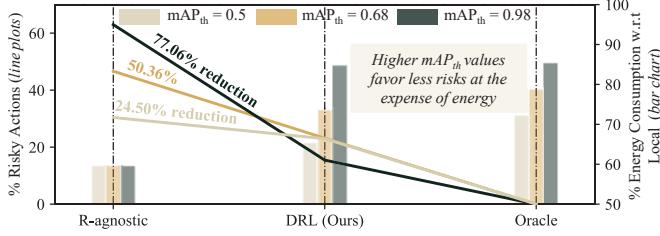


Figure 10: Comparing Energy and Risky Actions

policies, irrespective of the energy or robustness due to the L_{th} constraint. Such behavior is learned by our *DRL* solution given how the action selection frequencies closely imitate that of the *Oracle*.

Table 5: Energy analysis relative to local at $mAP_{th} = 0.98$

Metric	Local	R-agnostic	DRL	Oracle
Risky Actions (%)	0	63.37	14.54	0
Robust Actions (%)	100	36.63	85.46	100
Total Energy (kJ)	2.916	1.729	2.479	2.487
Total Energy Red. (%)	0	40.7	14.99	14.72

6.5 Energy Reduction vs Risky Actions

We also compare the energy savings relative to the pure local execution, *offload*, in addition to their risky behaviors. We first define *Risky Actions* as the fraction of offloading actions whose respective mAP scores fall below mAP_{th} , and *Robust Actions* as the fraction whose scores exceed the mAP_{th} . We compare the performance of each policy in Table 5, where although *R-agnostic* offers the highest energy reduction of 40.7% compared to *DRL*'s 14.99%, 63.37% of *R-agnostic*'s energy savings are *Risky Actions*, unlike *DRL* whose *Risky Actions* constitute 14.54% of the offloading decisions. Through extending this analysis further to entail multiple mAP_{th} values, i.e., a higher threshold means a stricter offloading constraint, we observe in Figure 10 that the robustness-aware *DRL* at higher mAP_{th} substantially reduces the percentage of risky offloads compared to the *R-agnostic* policy, with reductions of 24.50% and 77.06%, at $mAP_{th} = 0.50$ and $mAP_{th} = 0.98$ respectively.

7 CONCLUSION

In this work, we presented ROMANUS, a methodology for robust and efficient task offloading for multi-sensor autonomous driving systems (ADS). We first showed how to integrate optimal offloading points along the processing pipelines in a multi-sensor object detection module with late fusion, and then implemented a DRL-based runtime offloading that achieves 14.99% energy efficiency over pure local execution with up to 77.06% decrease in risky offloading actions from a robust-agnostic solution. This methodology can be generalized to a variety of sensors and fusion strategies depending on the underlying system structure and how robustness is characterized with regards to the application's primary task, which would form the basis for future research works along this direction.

REFERENCES

- [1] 2016. All new Teslas are equipped with NVIDIA's new Drive PX 2 AI platform for self-driving - Electrek. <https://electrek.co/2016/10/21/all-new-teslas-are-equipped-with-nvidias-new-drive-px-2-ai-platform-for-self-driving>.
- [2] 2017. Audi and NVIDIA team up to bring fully automated driving to the roads starting in 2020 accelerated with artificial intelligence. <https://media.audiusa.com/en-us/releases/104>.
- [3] 2021. Tesla Autopilot. <https://www.tesla.com/autopilot>
- [4] Sam Abuelsamid. 2020. Nvidia Cranks Up And Turns Down Its Drive AGX Orin Computers. *Forbes* (Jun 2020). <https://www.forbes.com/sites/samabuelsamid/2020/05/14/nvidia-cranks-up-and-turns-down-its-drive-agx-orin-computers>
- [5] Sabur Baidya, Yu-Jen Ku, Hengyu Zhao, Jishen Zhao, and Sujit Dey. 2020. Vehicular and edge computing for emerging connected and autonomous vehicle applications. In *ACM/IEEE Design Automation Conference (DAC'20)*.
- [6] Mingyue Cui et al. 2020. Offloading Autonomous Driving Services via Edge Computing. *IEEE Internet of Things Journal* 7, 10 (2020), 10535–10547.
- [7] Amir Erfan Eshratifar et al. [n.d.]. Bottlenet: A deep learning architecture for intelligent mobile cloud computing services. In *2019 IEEE/ACM Intl. Symp. on Low Power Electronics and Design (ISLPED'19)*.
- [8] Mark Everingham et al. 2010. The pascal visual object classes (VOC) challenge. *International journal of computer vision* 88, 2 (2010), 303–338.
- [9] Kaiming He et al. [n.d.]. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition (CVPR'2016)*.
- [10] Junxian Huang et al. 2012. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proc. of Intl. Conf. on Mobile Systems, Applications, and Services (MobiSys '12)*, 225–238.
- [11] Yiping Kang et al. [n.d.]. Neurosurgeon: Collaborative Intelligence Between the Cloud and Mobile Edge. In *Proc. of Intl. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS '17)*.
- [12] Guangli Li et al. 2018. Auto-tuning Neural Network Quantization Framework for Collaborative Inference Between the Cloud and Edge. In *Int. Conference on Artificial Neural Networks*, 402–411.
- [13] Shih-Chieh Lin et al. [n.d.]. The architectural implications of autonomous driving: Constraints and acceleration. In *Proc. of Intl. Conf. on Arch. Support for Programming Languages and Operating Systems (ASPLOS'18)*.
- [14] Shaoshan Liu, Liangkai Liu, Jie Tang, Bo Yu, Yifan Wang, and Weisong Shi. 2019. Edge computing for autonomous driving: Opportunities and challenges. *Proc. IEEE* 107, 8 (2019), 1697–1716.
- [15] Shaoshan Liu, Jie Tang, Zhe Zhang, and Jean-Luc Gaudiot. 2017. Computer architectures for autonomous driving. *Computer* 50, 8 (2017), 18–25.
- [16] Arnav Malawade et al. 2021. SAGE: A Split-Architecture Methodology for Efficient End-to-End Autonomous Vehicle Control. *ACM Transactions on Embedded Computing Systems (TECS)* 20, 5s (2021), 1–22.
- [17] Arnav Vaibhav Malawade et al. 2022. HydraFusion: Context-aware selective sensor fusion for robust and efficient autonomous vehicle perception. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPs)*.
- [18] Arnav Vaibhav Malawade, Trier Mortlock, and Mohammad Abdullah Al Faruque. 2022. EcoFusion: Energy-aware adaptive sensor fusion for efficient autonomous vehicle perception. *arXiv preprint arXiv:2202.11330* (2022).
- [19] Yoshitomo Matsubara et al. 2019. Distilled split deep neural networks for edge-assisted real-time systems. In *Proceedings of the 2019 Workshop on Hot Topics in Video Analytics and Intelligent Edges*.
- [20] Yoshitomo Matsubara et al. 2020. Head network distillation: Splitting distilled deep neural networks for resource-constrained edge computing systems. *IEEE Access* 8 (2020), 212177–212193.
- [21] Mohanad Odema et al. 2021. LENS: Layer distribution enabled neural architecture search in edge-cloud hierarchies. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 403–408.
- [22] Quazi Marufur Rahman, Niko Sünderhauf, and Feras Dayoub. 2021. Per-frame mAP Prediction for Continuous Performance Monitoring of Object Detection During Deployment. In *WACV (Workshops)*, 152–160.
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015), 91–99.
- [24] Babak Shahian Jahromi, Theja Tulabandhula, and Sabri Cetin. 2019. Real-time hybrid multi-sensor fusion framework for perception in autonomous vehicles. *Sensors* 19, 20 (2019), 4357.
- [25] Marcel Sheeny et al. 2021. RADIADE: A radar dataset for automotive perception in bad weather. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1–7.
- [26] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference*.
- [27] Haoxin Wang, BaekGyu Kim, Jiang Xie, and Zhu Han. 2019. E-auto: A communication scheme for connected vehicles with edge-assisted autonomous driving. In *2019 IEEE International Conference on Communications (ICC)*. IEEE.
- [28] Qiong Wu, Hanxu Liu, Ruhai Wang, Pingyi Fan, Qiang Fan, and Zhengquan Li. 2019. Delay-sensitive task offloading in the 802.11 p-based vehicular fog computing systems. *IEEE Internet of Things Journal* 7, 1 (2019), 773–785.
- [29] Danfei Xu et al. 2018. Pointfusion: Deep sensor fusion for 3D bounding box estimation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- [30] Jin Hyeok Yoo et al. [n.d.]. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *Computer Vision–ECCV 2020: Proceedings, Part XXVII 16*.
- [31] Ke Zhang et al. 2016. Energy-efficient offloading for mobile edge computing in 5G heterogeneous networks. *IEEE acc.* 4 (2016), 5896–5907.